arXiv:1703.07904v1 [stat.ME] 23 Mar 2017

Cross-Validation with Confidence

Jing Lei[1]

[1]Department of Statistics, Carnegie Mellon University

March 24, 2017

Abstract

Cross-validation is one of the most popular model selection methods in statistics and machine learning. Despite its wide applicability, traditional cross-validation methods tend to select overfitting models, unless the ratio between the training and testing sample sizes is much smaller than conventional choices. We argue that such an overfitting tendency of cross-validation is due to the ignorance of the uncertainty in the testing sample. Starting from this observation, we develop a new, statistically principled inference tool based on cross-validation that takes into account the uncertainty in the testing sample. This new method outputs a small set of highly competitive candidate models containing the best one with guaranteed probability. As a consequence, our method can achieve consistent variable selection in a classical linear regression setting, for which existing cross-validation methods require unconventional split ratios. We demonstrate the performance of the proposed method in several simulated and real data examples.

1 Introduction

Cross-validation [19, 1, 11] is one of the most popular methods for model and tuning parameter selection in statistics and machine learning, due to its conceptual simplicity and wide applicability. The basic idea of cross-validation is to fit and evaluate each candidate model on separate data sets so that the performance evaluation is unbiased. Despite its wide use, it is well-known that the traditional cross-validation methods, including the leave-one-out and V-fold variants, tend to select models that over fit. Early theoretical studies of cross-validation [18, 25, 24] indicate that, under a low-dimensional Gaussian linear model, cross-validation cannot consistently select the correct model unless the training-testing split ratio tends to zero, which precludes both the leave-one-out and V-fold cross-validation. Although these theory provide deep and important insight for cross-validation procedures, in practice it is very rare to use a small split ratio. In fact, most statistical packages use default split ratios such as nine-to-one (10-fold cross-validation) or four-to-one (5-fold

cross-validation). These conventional split ratios remain popular primarily because a small split ratio greatly reduces the training sample size and usually results in much less accurate model fitting.

In this paper we investigate the reason of the overfitting tendency of conventional cross-validation methods. It turns out that such an overfitting is due to the ignorance of uncertainty in the testing sample. In particular, overfitting occurs if a smaller average test error comes as a result of the random sampling of testing data instead of a truly superior fit.

To take into account the testing sample uncertainty and fix the issue of overfitting, we develop a hypothesis testing framework for cross-validation, which we call cross-validation with confidence (CVC). For each candidate model m, CVC tests the null hypothesis that the regression function estimated from model m has the smallest predictive risk among all fitted candidate models, and calculates a p-value by comparing the cross-validated residuals of *all* candidate models. The subset of candidate models for which the null hypotheses are not rejected is called "the admissible set" and can be viewed as a confidence set for model selection. We show that (i) if all the tests are conducted at type I error level α for some $\alpha \in (0, 1)$, then the admissible set contains the best fitted model with probability at least $1 - \alpha$, and (ii) with high probability, the admissible set only contains fitted models that are highly competitive. The main challenge of the testing problem is the high correlation and vastly different scaling between the cross-validated residuals. Our test uses some recent results in high-dimensional Gaussian comparison developed by [6, 5], which allows us to provide theoretical guarantees for our procedure under weak assumptions.

A classical and well-studied problem in the cross-validation literature is consistent variable selection in linear regression. Even in the low-dimensional case, cross-validation with a conventional split ratio is known to be inconsistent [18, 25, 24]. We show that the smallest model in the admissible set output by CVC achieves consistent variable selection even with a conventional split ratio.

Other related work. There are a few heuristic attempts to fix the overfitting issue of cross-validation. For example, [20] propose a heuristic bias correction for V-fold cross-validation, where the average cross-validated risk is offset by the standard deviation of the V individual subsample cross-validated risk estimates. [15] use a stability-based criterion, which is the ratio between the empirical variance of the fitted vectors and the squared ℓ_2-norm of the average fitted vector. Here the "empirical variance" and "average" are taken over repeated fitting using subsamples. Both methods still output a single candidate model. [13] describe a sample splitting method for assessing variable importance in building predictive regression models, which also uses a hypothesis testing idea. The theoretical properties are not clear for these methods. [23] propose a cross-validation based variable selection procedure which uses a three-way sample split and a least squares fit after Lasso.

[10] study model selection consistency of a version of cross-validation specifically designed for penalized likelihood estimators, where, in addition to assumptions on the signal pattern and the design matrix, the split ratio is required to vanish.

Another focus of the cross-validation literature is risk consistency. In an early work, [14] show that cross-validation can select a model with prediction risk close to the optimal candidate model. Similar results in high dimensional regression with Lasso have been established recently by [12, 7] for V-fold cross-validation procedures. [2] study the risk consistency of cross-validated Lasso with a specially designed 2-fold cross-validation variant. The additional value of our CVC method is that it outputs a few good candidate models with close-to-optimal performance, so that the user has more flexibility in choosing a particular model. For example, one can choose a model that is more parsimonious than the one given by standard cross-validation, or one can pick a model in the admissible set that contains a particularly desirable variable.

Model selection and tuning parameter selection. Depending on the context, model selection and tuning parameter selection may refer to different tasks. For example, if we have competing theories for the data generating mechanism, then we are choosing from a few candidate probabilistic models (such as linear versus nonlinear models). On the other hand, tuning parameter selection often concerns choosing the right amount of regularization for a class of estimators. In many cases a particular value of tuning parameter implies searching over a restricted subset in the parameter space, which can be viewed as a model in a broad sense. Therefore, in this paper we will use the unified terminology "model selection", even though our discussion also includes tuning parameter selection.

Organization of the paper. In Section 2 we briefly review cross-validation and investigate the cause of overfitting. In Section 3 we present the CVC method and its extensions. In Section 4 we provide general coverage and performance guarantees of the admissible set, and prove variable selection consistency of the CVC method in a classical linear regression setting. Section 5 presents simulated and real data examples. Section 6 concludes the paper with some implementation details and open problems. Appendix A contains technical proofs of main results.

2 Review of cross-validation

Suppose we have data $D = (X_i, Y_i)_{i=1}^n$ independently drawn from a common distribution P on $\mathbb{R}^p \times \mathbb{R}$ satisfying

$$Y_i = f(X_i) + \epsilon_i$$

where $f : \mathbb{R}^p \mapsto \mathbb{R}$ is an unknown function and ϵ satisfies $\mathbb{E}(\epsilon \mid X) = 0$. We are interested in finding an approximation of f so that we can predict the values of Y for future observations of X. Let $\hat{f}$ be an estimate of f. We evaluate the quality of $\hat{f}$ by the predictive risk

$$Q(\hat{f}) = \mathbb{E}\left[\ell(\hat{f}(X), Y) \mid \hat{f}\right] ,$$

where (X, Y) is a future random draw from P, and $\ell(\cdot, \cdot)$ is a loss function. A typical example of $\ell(\cdot, \cdot)$ is the squared loss $\ell(\hat{y}, y) = (\hat{y} - y)^2$.

2.1 Cross-validation.

Given a finite set of candidate models $\mathcal{M} = \{1, 2, ..., M\}$, for each candidate model m one can estimate f by $\hat{f}_m$ using an estimating procedure specified by m. Here m can be interpreted as the index of a competing theory about f, or the index of a tuning parameter value.

The simplest form of cross-validation, which we call single split cross-validation, works as follows. First we split the index set $\{1, ..., n\}$ randomly into a training set I_{tr} and a testing set $I_{\text{te}} = I_{\text{tr}}^c$. The data is correspondingly split into a training data $D_{\text{tr}} = \{(X_i, Y_i) : i \in I_{\text{tr}}\}$ and a testing data $D_{\text{te}} = \{(X_i, Y_i) : i \in I_{\text{te}}\}$. For each candidate model m, we obtain an estimate $\hat{f}_m$ using D_{tr}. The quality of $\hat{f}_m$ is approximated using the empirical average loss on D_{te}:

$$\widehat{Q}(\hat{f}_m) = \frac{1}{n_{\text{te}}} \sum_{i \in I_{\text{te}}} \ell(\hat{f}_m(X_i), Y_i) ,$$

where n_{te} is the size of D_{te}. Finally cross-validation chooses the model that gives the smallest empirical risk on testing data:

$$\hat{m}_{\text{cv}} = \arg\min_m \widehat{Q}(\hat{f}_m) . \tag{1}$$

In practice, cross-validation often uses a V-fold split, and repeats the calculation of $\hat{f}_m$ and $\widehat{Q}(\hat{f}_m)$ by rotating each fold as the testing data. The final approximated predictive risk is the average value of $\widehat{Q}(\hat{f}_m)$ over the V folds. The extreme case of $V = n$ is known as the leave-one-out cross-validation. In modern applications it is often expensive and unnecessary to re-fit the model n times, so most statistical packages use 10-fold or 5-fold cross-validation as default choices.

2.2 What does cross-validation try to optimize?

Conditioning on the sample split and the training data D_{tr}, the Law of Large Numbers implies that $\widehat{Q}(\hat{f}_m)$ is close to $Q(\hat{f}_m)$. Therefore, cross-validation approximately minimizes the predictive risk over all candidate estimates $\{\hat{f}_m : m \in \mathcal{M}\}$.

There are two levels of inference we can hope for in the context of cross-validation.

1. Conditional inference. The task is to find the best estimate of f from $\{\hat{f}_m : m \in \mathcal{M}\}$, where $\hat{f}_m$ is obtained from the training data. Such an inference is conditional on D_{tr} and only concerns the randomness in D_{te}.

2. Unconditional inference. The task is to find the model m that produces the best estimate of f in a certain probabilistic sense. To this end, some additional assumptions on the stability and consistency of the estimate are necessary, because we need to make sure that the estimate $\hat{f}_m$ obtained from D_{tr} is representative for the overall performance of model m. Such assumptions are explicitly stated in [24], and implicitly established in the proof in [18, 25] under the context of low-dimensional linear regression.

In practice, we often want an estimate of f in addition to choosing a model. A common practice is to estimate f using the whole data set with $\hat{m}_{\text{cv}}$. This would be reasonable only if the size of D_{tr} is close to the whole sample size, because in many problems the optimal choice of tuning parameter depends on the sample size. Moreover, given a finite sample, an overly small training sample size will lead to noisy estimates and hence less accurate model selection. Throughout this paper we will focus on common choices of split ratio, such as 5-fold or 10-fold cross-validation.

2.3 When and why does cross-validation overfit?

We start our discussion of overfitting from a simple example.

Example 1. *Let $Y = \mu + \epsilon$ where $\epsilon \sim N(0, 1)$. The covariate X is irrelevant here. Let $\mathcal{M} = \{0, 1\}$ where $m = 0$ corresponds to the model $\mu = 0$ so $\hat{f}_0 \equiv 0$, and $m = 1$ corresponds to the model $\mu \in \mathbb{R}$ so $\hat{f}_1(x) \equiv \hat{\mu}_{\text{tr}}$ with $\hat{\mu}_{\text{tr}}$ being the training sample mean of Y. We consider the squared error loss. Now assume $\mu = 0$, then $\hat{f}_0$ is always a better estimate than $\hat{f}_1$. But*

$$\widehat{Q}(\hat{f}_1) - \widehat{Q}(\hat{f}_0) = \hat{\mu}_{\text{tr}}^2 - 2\hat{\mu}_{\text{tr}}\bar{\epsilon}_{\text{te}}\,,$$

where $\bar{\epsilon}_{\text{te}}$ is the testing data sample mean of ϵ. Because $\sqrt{n_{\text{tr}}}\hat{\mu}_{\text{tr}}$ and $\sqrt{n_{\text{te}}}\bar{\epsilon}_{\text{te}}$ are independent standard normal random variables, there is positive probability that $\widehat{Q}(\hat{f}_1) - \widehat{Q}(\hat{f}_0) < 0$. This probability becomes larger when n_{te} is smaller.

In Example 1, cross-validation makes decision based on the sign of $\widehat{Q}(\hat{f}_1) - \widehat{Q}(\hat{f}_0)$, which contains a true signal part $\hat{\mu}_{\text{tr}}^2$ and a noise part $-2\hat{\mu}_{\text{tr}}\bar{\epsilon}_{\text{te}}$. Overfitting happens if the noise overrides the signal. In other words, cross-validation overfits because there is a very competitive overfitting model and the procedure cannot tell if a particular sign of a sample average is due to a true signal or just randomness of the testing data.

3 Cross-validation with confidence

Given the fitted regression functions $\{\hat{f}_m : m \in \mathcal{M}\}$ from the training data, we propose a new procedure for finding good candidate estimates. The procedure, called cross-validation with confidence (CVC), fully takes into account the uncertainty in $\widehat{Q}(\hat{f}_m)$ due to the random sampling of testing data using a statistical hypothesis testing framework. The output of CVC is a subset $\mathcal{A}_{\text{cvc}} \subseteq \{\hat{f}_m : m \in \mathcal{M}\}$ of candidate estimates that have highly competitive predictive performance.

As explained in Section 2, cross-validation overfits when it fails to calibrate the uncertainty of testing data. Here we focus on the conditional inference given the training data D_{tr} and the candidate estimates $\{\hat{f}_m : m \in \mathcal{M}\}$. Recall that we evaluate a fitted model $\hat{f}$ by $Q(\hat{f}) = \mathbb{E}(\ell(\hat{f}(X), Y) \mid \hat{f})$. For each $m \in \mathcal{M}$, we consider a hypothesis testing problem

$$\begin{aligned} H_{0,m}: \;& Q(\hat{f}_m) \le Q(\hat{f}_{m'}) \;\; \forall\, m' \neq m \\ \text{vs} \\ H_{1,m}: \;& Q(\hat{f}_m) > Q(\hat{f}_{m'}) \;\; \text{for some} \;\; m' \neq m\,. \end{aligned} \tag{2}$$

The hypothesis $H_{0,m}$ means that the fitted model $\hat{f}_m$ has the best predictive risk among all fitted models, and $H_{1,m}$ means that there exists another fitted model whose predictive risk is strictly less than $\hat{f}_m$.

Let $\alpha \in (0,1)$ be a pre-chosen level of type I error. If we have obtained a p-value, denoted by $\hat{p}_m$, for the testing problem (2) for each m, the CVC procedure outputs the *admissible set*

$$\mathcal{A}_{\text{cvc}} = \{\hat{f}_m : m \in \mathcal{M},\; \hat{p}_m > \alpha\}, \tag{3}$$

which contains all the fitted models for which $H_{0,m}$ is not rejected.

3.1 Calculating the p-values.

Fix an $m \in \mathcal{M}$, and define random vector $\xi_m = (\xi_{m,j} : j \neq m)$ as

$$\xi_{m,j} = \ell(\hat{f}_m(X), Y) - \ell(\hat{f}_j(X), Y), \tag{4}$$

where (X, Y) is a fresh independent sample from the same joint distribution. Let $\mu_{m,j} = \mathbb{E}(\xi_{m,j} \mid \hat{f}_m, \hat{f}_j)$. Then the hypotheses in (2) can be written equivalently as

$$H_{0,m} : \max_{j \neq m} \mu_{m,j} \le 0, \quad \text{vs} \quad H_{1,m} : \max_{j \neq m} \mu_{m,j} > 0\,. \tag{5}$$

(5) is a multivariate mean testing problem. There are a few challenges. First, the dimensionality can be high, as the number of candidate models or tuning parameter values can be

on the order of hundreds. Second, the coordinates of ξ_m are often highly correlated, since they are calculated from prediction errors of similar models fitted from the same training data and evaluated on the same testing data. Third, the coordinates of ξ_m usually have vastly different scaling, depending on the similarity between model j and model m. For example, if $\hat{f}_m$ and $\hat{f}_j$ are both $\sqrt{n}$-consistent then $\xi_{m,j}$ has a variance of order $O(n^{-1})$. But if $\hat{f}_m$ is $\sqrt{n}$-consistent but $\hat{f}_j$ is seriously underfitting (e.g., a linear model missing a relevant variable, or Lasso with a very large λ), then the variance of $\xi_{m,j}$ can be as large as a constant.

Considering all these challenges, we derive our test statistic and p-value using a recent result of Gaussian comparison and bootstrap of high-dimensional sample means due to [4], which provides an accurate approximation of the maximum of a high-dimensional sample mean using the maximum of a corresponding Gaussian random vector.

The p-value $\hat{p}_m$ is calculated as follows, using a studentized Gaussian multiplier bootstrap approach.

1. For $i \in I_{\text{te}}$, let $\xi_m^{(i)}$ be the realization of ξ_m on the testing data point (X_i, Y_i):
$$\xi_m^{(i)} = (\xi_{m,j}^{(i)} : j \neq m) \quad \text{where} \quad \xi_{m,j}^{(i)} = \ell(\hat{f}_m(X_i), Y_i) - \ell(\hat{f}_j(X_i), Y_i)\,.$$
2. For each $j \neq m$, let $\hat{\mu}_{m,j}$ and $\hat{\sigma}_{m,j}$ be the sample mean and standard deviation of $\{\xi_{m,j}^{(i)} : i \in I_{\text{te}}\}$.
3. Let $T = \max_{j \neq m} \sqrt{n_{\text{te}}} \frac{\hat{\mu}_{m,j}}{\hat{\sigma}_{m,j}}$.
4. For $b = 1, ..., B$
   (a) Generate iid standard Gaussian random variables ζ_i, $i \in I_{\text{te}}$.
   (b) Let
   $$T_b^* = \max_{j \neq m} \frac{1}{\sqrt{n_{\text{te}}}} \sum_{i \in I_{\text{te}}} \frac{\xi_{m,j}^{(i)} - \hat{\mu}_{m,j}}{\hat{\sigma}_{m,j}} \zeta_i.$$
5. $\hat{p}_m = B^{-1} \sum_{b=1}^{B} \mathbf{1}(T_b^* > T)$.

The Gaussian comparison and bootstrap results automatically take into account of the dimensionality and correlation between coordinates of ξ_m. The studentization addresses the issue of scaling difference. A similar studentized Gaussian approximation has been considered in an unpublished manuscript [5] for different applications.

The way we compute $\hat{p}_m$ and the definition of $\mathcal{A}_{\text{cvc}}$ guarantee that the latter is non-empty, and also reveal the relationship between standard cross-validation and CVC.

Proposition 3.1. *If $\alpha < 0.5$, then*

$$\lim_{B\to\infty} \mathbb{P}\left(\hat{m}_{\mathrm{cv}} \in \mathcal{A}_{\mathrm{cvc}}\right) = 1\,,$$

where $\hat{m}_{\mathrm{cv}}$ is the standard cross-validation output given in (1).

Proof. Be definition of $\hat{m}_{\mathrm{cv}}$, we always have $T \le 0$. Because $\alpha < 0.5$, the upper α quantile of the maximum of a zero-mean Gaussian vector (T_b^*) must be positive. The claim then follows from the Weak Law of Large Numbers. ∎

Following the same reasoning, we can see that among all $m \in \mathcal{M}$, the test statistic T is only non-positive when m is the cross-validation choice $\hat{m}_{\mathrm{cv}}$. Thus, empirically $\hat{m}_{\mathrm{cv}}$ almost always gives the largest CVC p-value among all candidate models.

3.2 V-fold CVC.

Like the standard cross-validation, CVC can be implemented in the popular and more stable V-fold manner. For a positive integer $V \ge 2$ (for example, $V = 5, 10$), assume that the index set $\{1, 2, ..., n\}$ is randomly split into V equal-sized folds $(I_v : 1 \le v \le V)$, with a corresponding data set split $D_v = \{(X_i, Y_i) : i \in I_v\}$ for $1 \le v \le V$. Let v_i be the fold that contains sample point i. The V-fold CVC proceeds the in same way as the V-fold cross-validation to compute the prediction error $\ell(\hat{f}_{m,v_i}(X_i), Y_i)$ for all $m \in \mathcal{M}$ and $1 \le i \le n$, where $\hat{f}_{m,v}$ is the regression function estimated from model m on all data except the vth fold.

In the next step, unlike the standard V-fold cross-validation which simply averages $\ell(\hat{f}_{m,v_i}(X_i), Y_i)$ over all $i \in \{1, ..., n\}$ as the estimated predictive risk of model m, V-fold CVC computes the p-value $\hat{p}_m$ following steps (1)-(5) in Section 3.1 using the $n \times (M-1)$ data matrix

$$\left[\xi_{m,j}^{(i)}\right]_{1\le i\le n, j\ne m}, \quad \text{where} \quad \xi_{m,j}^{(i)} = \ell(\hat{f}_{m,v_i}(X_i), Y_i) - \ell(\hat{f}_{j,v_i}(X_i), Y_i)\,.$$

Finally, we remark that Proposition 3.1 also holds for V-fold CVC.

4 Theoretical Properties

In our theoretical analysis of CVC, we first consider conditional inference, which is concerned with finding good candidate estimates $\{\hat{f}_m : m \in \mathcal{M}\}$ given the training data D_{tr}. Next we consider unconditional inference, which aims at finding the best m that corresponds to the "true model", under the context of subset selection for classical linear models.

Our arguments require the estimated prediction errors realized on the testing data to be independent conditioning on the training data, which is the case in the single split cross-validation. Some heuristic discussion about the validity of V-fold CVC is presented at the end of Section 4.1 after the statement of the main result.

Our technical conditions and theorems are stated in an asymptotic manner so that the sample size n grows to infinity, while other quantities, such as the number of models M and the tail behavior of the prediction error, are allowed to change with n. Also we always assume that the test sample size is proportional to the total sample size: $n_{\text{te}}/n \in [\delta, 1-\delta]$ for some $\delta \in (0, 1/2)$.

4.1 General quality guarantees of $\mathcal{A}_{\text{cvc}}$.

Our main technical assumption is the tail behavior of the difference of prediction loss $\xi_{m,j} = \ell(\hat{f}_m(X), Y) - \ell(\hat{f}_j(X), Y)$, where the randomness is in the future data pair (X, Y). We assume sub-exponential tail probabilities for standardized $\xi_{m,j}$, characterized by the Orlicz ψ_1 norm. Given $a \geq 1$, the Orlicz ψ_a norm of a random variable Z is

$$\|Z\|_{\psi_a} = \inf\left\{C > 0 : \mathbb{E}\exp\left(\frac{|Z|}{C}\right)^a \leq 2\right\}.$$

Having a finite Orlicz ψ_1 norm is equivalent to having a sub-exponential tail, and implies finite moments of all orders and exponential concentration of sample mean. A finite Orlicz ψ_2 norm is equivalent to a sub-Gaussian tail. See [22] and [21] for further discussion on the Orlicz norm.

Following common practice, our analysis of the multiplier bootstrap assumes that B is large enough so the bootstrap sampling variability is negligible. That is, we take $\hat{p}_m$ as the limiting value when $B \to \infty$.

Theorem 4.1. *Let $\mu_{m,j} = \mathbb{E}(\xi_{m,j} \mid \hat{f}_m, \hat{f}_j)$, $\sigma^2_{m,j} = \mathrm{Var}(\xi_{m,j} \mid \hat{f}_m, \hat{f}_j)$. Assume*

$$\sup_{m \neq j} \left\| \frac{\xi_{m,j} - \mu_{m,j}}{\sigma_{m,j}} \right\|_{\psi_1} \leq A_n, \tag{6}$$

for all $m \neq j$ with

$$\frac{(A_n \vee 1)^6 \log^7(M \vee n)}{n^{1-c}} \to 0, \tag{7}$$

for some $c \in (0, 1)$. The following hold for the CVC procedure.

1. *If $\max_{j \neq m} \frac{\mu_{m,j}}{\sigma_{m,j}} \leq r_n \sqrt{\frac{1}{n \log(M \vee n)}}$ for some $r_n = o(1)$, then $\mathbb{P}(m \in \mathcal{A}_{\text{cvc}}) \geq 1 - \alpha + o(1)$.*
2. *If $\alpha \in (n^{-1}, 1)$ and $\max_{j \neq m} \frac{\mu_{m,j}}{\sigma_{m,j}} \geq c'(A_n \vee 1)\sqrt{\frac{\log(M \vee n)}{n}}$ for a constant c' depending only on n/n_{te}, then $\mathbb{P}(m \in \mathcal{A}_{\text{cvc}}) = o(1)$.*

We prove Theorem 4.1 in Appendix A.1. Part (1) of Theorem 4.1 guarantees the inclusion of competitive estimates. If $\hat{f}_m$ is the best estimate in $\{\hat{f}_m : m \in \mathcal{M}\}$, then $H_{0,m}$ is true and m will be included in $\mathcal{A}_{\text{cvc}}$ with probability at least $1 - \alpha + o(1)$ by just taking $r_n = 0$ in part (1). In fact, the proof reveals that in the trivial case where $\mu_{m,j} = \mathbb{E}(\xi_{m,j} \mid \hat{f}_m, \hat{f}_j) = 0$ for all $j \neq m$, we have $\mathbb{P}(m \in \mathcal{A}_{\text{cvc}}) = 1 - \alpha + o(1)$.

Remark 1. *Theorem 4.1 holds even when α vanishes as n goes to infinity. We will use this result to prove variable selection consistency of CVC for low-dimensional linear regression in the next subsection.*

Remark 2. *The sub-exponential tail condition (6) can be satisfied, for example, for linear regression with squared loss if the covariate X and noise ϵ are sub-Gaussian. The Orlicz ψ_1 norm in (6) will depend on the estimated coefficients $\hat{\beta}_m$, which are typically restricted in a bounded set with high probability on D_{tr}. The proofs of our main results in the next subsection indeed provide further details on establishing (6) in least square linear regression assuming sub-Gaussian covariates and noise.*

In V-fold CVC, due to the dependence between the $\xi_{m,j}^{(i)}$'s among different folds, it is hard to rigorously guarantee the performance. Empirically, we found that V-fold CVC has much better power in eliminating suboptimal estimates. We provide a heuristic explanation below.

Suppose there are only two candidate models $\mathcal{M} = \{0, 1\}$ and that $m = 0$ tends to give better estimates (see Example 1 for a concrete example). Let $\hat{\mu}_{1,0,v}$ and $\hat{\sigma}^2_{1,0,v}$ be the estimated mean and variance of $\xi_{1,0}$ using only fold v as test data. Let $\hat{\mu} = V^{-1} \sum_{v=1}^{V} \hat{\mu}_{1,0,v}$ and $\hat{\sigma}^2 = V^{-1} \sum_{v=1}^{V} (\hat{\sigma}^2_{1,0,v} + \hat{\mu}^2_{1,0,v}) - \hat{\mu}^2$ be the overall sample mean and variance of the aggregated data matrix used in CVC. Assuming that the estimators are stable and the variation of estimates across different folds is negligible, we have $\hat{\mu}_{1,0,v} \approx \hat{\mu} > 0$ and $\hat{\sigma} \approx \hat{\sigma}_{1,0,v}$ for all v. Thus, the signal strength is amplified by a factor of $V^{1/2}$ in V-fold CVC compared to using a single fold as testing data.

4.2 Model selection consistency in classical linear models.

Now we show that CVC can be used to overcome the overfitting issue of standard cross-validation in a classical linear regression setting. Assume that the regression function is linear

$$Y = X^T \beta + \epsilon \,, \tag{8}$$

where $X \in \mathbb{R}^p$ has covariance Σ, and ϵ is independent noise with mean zero and variance σ^2. Here for brevity we assume that p is fixed. The argument and results can be extended to the case that p grows as a small polynomial of n using a union bound argument.

Given a collection of subsets $\mathcal{J} = \{J_1, ..., J_M\} \subseteq 2^p$, we would like to find the m^* such that $J_{m^*} = \{j : \beta_j \neq 0\}$, assuming that the true model is included in the candidate set. The standard cross-validation (using a single split) estimates $\hat{\beta}_m$ by applying a least squares fit using the training data and covariates in J_m, and evaluates the model by the residual sum of squares on the testing data. CVC works analogously with the same least squares fitting and uses the squared residual as loss function. For any given α, CVC outputs an admissible set of candidate models. To select the correct model, we consider the most parsimonious admissible model given by CVC

$$\hat{m}_{\text{cvc}} = \arg\min_{m \in \mathcal{A}} |J_m|. \tag{9}$$

According to Proposition 3.1, $\mathcal{A}_{\text{cvc}}$ is always non-empty as long as $\alpha < 0.5$. So $\hat{m}_{\text{cvc}}$ is always well-defined when $\alpha < 0.5$.

We consider the following conditions.

(A1) $X = \Sigma^{1/2}\tilde{X}$, where $\tilde{X}$ has identity covariance and satisfies $\|\tilde{X}\|_{\psi_2} \equiv \sup_{v:\|v\|_2=1} \|v^T\tilde{X}\|_{\psi_2} < \infty$.

(A2) $\lambda_{\min} > 0$, where $\lambda_{\min}$ is the minimum eigenvalue of Σ.

(A3) $\|\epsilon\|_{\psi_2} < \infty$.

These are rather mild conditions for linear regression problems. (A1) and (A3) assumes that both X and ϵ are sub-Gaussian, so that the squared residual is sub-exponential. (A2) is necessary for the best model to be uniquely defined.

Theorem 4.2. *Assume model (8) and assumptions A1-A3 hold. Let $\hat{m}_{\text{cvc}}$ be given as in (9) with $\mathcal{A}$ being the CVC admissible set using type I error level $\alpha_n \to 0$ and $\alpha_n > n^{-1}$. Then $\mathbb{P}(\hat{m}_{\text{cvc}} = m^*) \to 1$.*

We prove Theorem 4.2 in Appendix A.2. One of the main challenges is to carefully calibrate the mean, variance, and Orlicz ψ_1 norm of $\xi_{m,m'}$, especially when both J_m and $J_{m'}$ contains all the relevant variables. Another challenge is to show that $\hat{f}_{m^*}$ has the smallest risk with high probability, which involves the Kolmogorov-Rogozin anti-concentration inequality [17] since we do not assume Gaussianity for X or ϵ.

5 Numerical experiments

We illustrate the performance of the CVC method under several simulation settings and a real data example.

5.1 Simulation 1: Coverage and power of CVC.

We use this simulation to demonstrate the ability of CVC to correctly cover the best fitted model with desired probability, as well as to effectively eliminate clearly inferior models (Theorem 4.1). We generate data from a linear regression model

$$Y = X^T\beta + \epsilon, \tag{10}$$

where $X \in \mathbb{R}^{200}$ has multivariate Gaussian distribution $N(0, \Sigma)$, and ϵ is independent standard Gaussian noise. We consider the squared loss. For any estimate $\hat{\beta}$ of β, the model (10) leads to a closed-form formula for $Q(\hat{\beta})$:

$$Q(\hat{\beta}) = \mathbb{E}\left[(Y - X^T\hat{\beta})^2 \mid \hat{\beta}\right] = (\hat{\beta} - \beta)^T\Sigma(\hat{\beta} - \beta) + 1\,.$$

We consider two settings of Σ. In the *identity setting*, Σ is the 200×200 identity matrix. In the *correlated setting*, the diagonal entries of Σ are 1 and off-diagonal entries are 0.5. We also consider two settings of β. In the *simple setting*, $\beta = (1, 1, 1, 0, ..., 0)^T$ (the first three entries of β equal 1 and the rest are 0). In the *mixed setting*, $\beta = (1, 1, 1, 0.7, 0.5, 0.3, 0, ..., 0)^T$ (the first three entries of β are 1, followed by 0.7, 0.5, 0.3 in the next three entries, and 0 elsewhere).

We consider tuning parameter selection for two popular sparse regression methods: the Lasso and forward stepwise regression. For the Lasso, we first obtain a sequence of 50 tuning parameters using `R` package `glmnet` on the whole data set. Then proceed with 5-fold CVC with $\alpha = 0.05$. Because cross-validation (and hence CVC) does not automatically output a final estimate of the regression coefficient, we evaluate a tuning parameter value by the prediction risk of the fit obtained from the whole data. As mentioned before, the best Lasso tuning parameter for sample size $n = 160$ (the training sample size in 5-fold cross-validation) may differ from that for sample size $n = 200$. Thus we scale the tuning parameter by a factor of $\sqrt{4/5}$ when calculating the final estimate from the whole data set. For forward stepwise regression, we use the `lars` package with maximum number of steps being 20.

We run the simulation with sample size $n = 200$, and repeat 100 times for each fitting method and each combination of β and Σ. The results are summarized in Table 1, which reports the proportion of time the best fitted model is covered in $\mathcal{A}_{\text{cvc}}$ ("coverage"), and the average size of $\mathcal{A}_{\text{cvc}}$ ("$|\mathcal{A}_{\text{cvc}}|$"), over the 100 repetitions. The standard errors of the mean values are given in the parentheses. To compare with the standard 5-fold cross-validation, we also report the proportion of time that the best tuning parameter value is chosen by standard cross-validation ("cv opt").

It is clear that $\mathcal{A}_{\text{cvc}}$ covers the best fitted model with desired probability. Also the standard cross-validation often misses the best model, especially in the case of Lasso, where the tuning parameter changes continuously.

setting of (Σ, β)	lasso			forward		
	coverage	$\lvert\mathcal{A}_{\text{cvc}}\rvert$	cv opt.	coverage	$\lvert\mathcal{A}_{\text{cvc}}\rvert$	cv opt.
identity, simple	.92 (.03)	5.1 (.19)	.27 (.04)	1 (0)	3.7 (.29)	.87 (.03)
identity, mixed	.95 (.02)	5.1 (.18)	.37 (.05)	.95 (.02)	5.2 (.33)	.58 (.05)
correlated, simple	.96 (.02)	7.5 (.18)	.18 (.04)	.97 (.02)	4.1 (.31)	.80 (.04)
correlated, mixed	.93 (.03)	7.4 (.23)	.19 (.04)	.93 (.03)	6.3 (.36)	.44 (.05)

Table 1: Results for Simulation 1. Reported are proportion of simulation repetitions in which $\mathcal{A}_{\text{cvc}}$ covers the best fitted candidate model ("coverage"), and the average size of $\mathcal{A}_{\text{cvc}}$ ("$|\mathcal{A}_{\text{cvc}}|$"). Standard errors are included in the parentheses. The columns "cv opt." reports the proportion of simulation repetitions in which the best fitted model is chosen by standard 5-fold cross-validation.

5.2 Simulation 2: subset selection consistency in linear models.

In this simulation we demonstrate the model selection consistency theory for CVC developed in Section 4.2. We adopt the synthetic data set in [18], where the covariate X has five coordinates including intercept, with a sample size $n = 40$ (see Table 1 of [18] for the complete data set). The response variable Y is generated by the linear model (10) with a pre-determined β and independent standard Gaussian noise ϵ. We use $\beta = (2, 9, 0, 4, 8)^T$, and a set of candidate models $\mathcal{M} = \{(1, 4, 5), (1, 2, 4, 5), (1, 3, 4, 5), (1, 2, 3, 4, 5)\}$, which contains the true model ($m^* = 2$) as well as both overfitting and underfitting candidate models.

The sample size $n = 40$ is far from the regime of model selection consistency. To illustrate the effect of sample size, we generate additional sample points with X drawn from a multivariate Gaussian distribution whose parameters are given by the sample mean and covariance of the original data set. We consider varying sample sizes between $n = 40$ and $n = 160$, where $n = 40$ corresponds to the original data set. We implement 5-fold CVC with $\alpha = 0.05$ and select the model according to (9). The selected model is denoted as "cvc.min". The performance is compared with the standard 5-fold cross-validation.

Figure 1 summarizes the proportion of time each method correctly selects the model, over 1000 independently generated data and independent sample splitting. For a very small sample size $n = 40$, the CVC method does not perform as well as the standard cross-validation, as it often selects the underfitting model $(1, 4, 5)$. The performance of CVC increases rapidly as the sample size n increases, with a nearly perfect rate of correct selection as soon as n reaches 80. Meanwhile, the performance of the standard cross-validation does not improve as n increases, with rates of correct selection staying around 0.8. This agrees with the theory that CVC can consistently select the true model, while the standard cross-validation tends to overfit as long as the training sample ratio stays constant.

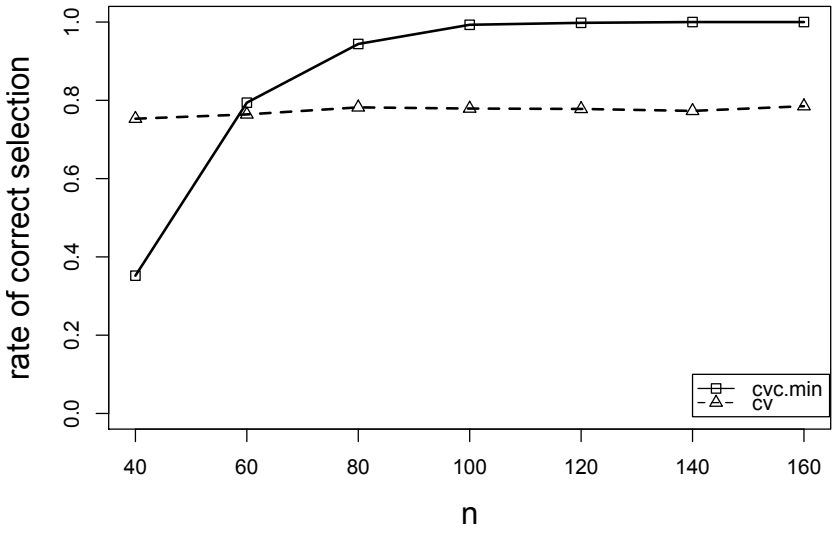

Figure 1: Results for Simulation 2. Reported are rates of correct model selection over 1000 repetitions for 5-fold CVC and standard 5-fold cross-validation for the synthetic data with sample sizes varying from 40 to 160.

5.3 Simulation 3: high-dimensional variable selection.

Section 5.1 demonstrates the validity of CVC in terms of prediction risk. Now we investigate the variable selection performance of CVC using forward stepwise selection and the Lasso in a moderately high dimensional setting. We use the same settings as in Simulation 1, except that we only consider the case of a simple β vector. When β has small non-zero coordinates such as in the mixed setting, model selection is hard anyway. The Lasso estimator is implemented with a least squares re-fitting step. That is, for each value of tuning parameter, we first fit an ordinary Lasso estimator, and then re-fit a least squares estimator using the variables selected by the Lasso. This is to reduce the bias in the original Lasso estimator, so that model selection and prediction risk minimization become better aligned. Similar as in Simulation 1, we use a sample size $n = 200$ across all settings, and repeat each setting 100 times. We implement 5-fold CVC model selection according to (9) with $\alpha = 0.05$, and compare with the standard 5-fold cross-validation (denoted as "cvc.min" and "cv", respectively).

Table 2 summarizes the results. Here we include an "oracle" method, which uses the tuning parameter value that gives the smallest predictive risk. We see that in the simple case, where the covariance of X is identity, the CVC method gives perfect model selection, while the standard cross-validation method makes a mistake for more than 10% of the time. In the harder case where the coordinates of X are highly correlated, the CVC implemented

setting of (Σ, β)	forward			lslasso		
	oracle	cvc.min	cv	oracle	cvc.min	cv
identity, simple	1	1	.87	1	1	.88
correlated, simple	1	.97	.80	.87	.85	.71

Table 2: Results for Simulation 3. Reported are proportion of the 100 simulated data sets for which the correct model is chosen. The "oracle" method corresponds to the tuning parameter value that gives the smallest prediction risk. Both CVC and cross-validation are implemented with 5-fold.

with forward selection still has nearly perfect rate of correct selection. For the Lasso with least squares, model selection is hard even for the oracle method when the covariates are highly correlated. Here CVC performs similarly to the oracle, with a 85% rate of correct selection, while the standard cross-validation has a 71% rate of correct selection.

5.4 The diabetes data example.

We apply the CVC method to the diabetes data, which has been used in [9] to illustrate the LARS algorithm in sparse linear regression. The data contains $n = 442$ sample points, with ten covariates including age, sex, body mass index, blood pressure, and six blood serum measurements. If we consider all the quadratic terms as in the literature, there will be 64 predictors (the sex variable is coded as $\{0, 1\}$ and has no square term). The response variable is a quantitative measure of diabetes progression one year after the covariates are measured. Following common practice, we center and standardize all variables.

Figure 2 shows the result of 5-fold CVC with $\alpha = 0.05$ on the full data set. The base estimator is the Lasso, with 50 equally spaced values (on the log scale) of the penalty parameter λ. We plot the cross-validated test errors for all candidate values of λ. The triangle points correspond to the tuning parameter values in $\mathcal{A}_{\text{cvc}}$, and the solid triangle corresponds to the λ chosen by standard 5-fold cross-validation. 5-fold CVC outputs five candidate values of λ, one of which gives a very close test error as the standard cross-validation but a more parsimonious model fit.

To investigate the predictive performance of the final estimate, we split the data so that 300 sample points are used for estimation, including implementing cross-validation or CVC, and the remaining 142 sample points are held out for test error calculation. To be specific, we use CVC and standard cross-validation, both 5-fold, to pick λ, and then obtain the final estimate of regression coefficient using Lasso with the chosen λ on the full training sample of 300 data points. Then we use the untouched set of 142 hold-out sample points to calculate the test error for the final estimate. We repeat this experiment 100 times, each time with a new hold-out split.

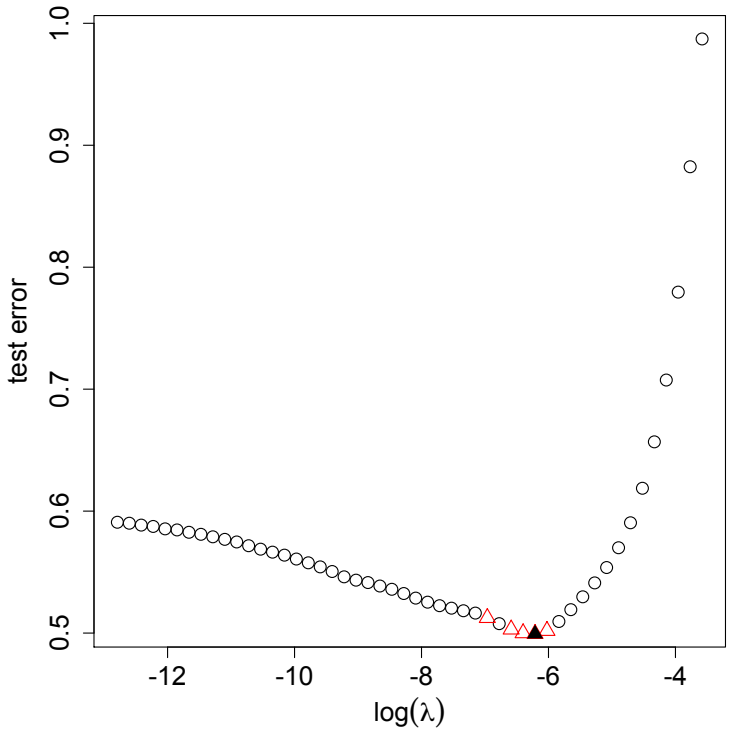

Figure 2: Diabetes data example. The cross-validation test errors for 50 values of λ. The triangles correspond to the λ values in $\mathcal{A}_{\text{cvc}}$ given by 5-fold CVC with $\alpha = 0.05$. The solid triangle corresponds to the λ chosen by standard 5-fold cross-validation.

The results are summarized in Figure 3. For CVC, we consider the most parsimonious model in $\mathcal{A}_{\text{cvc}}$ ("cvc.min"). We see that the most parsimonious CVC model gives very close test error as the standard cross-validation, but has much fewer predictors.

It is also remarkable that in 87 out of 100 repetitions, the set $\mathcal{A}_{\text{cvc}}$ contains a model that gives strictly smaller test error on the hold-out data than the standard cross-validation. In the remaining 13 repetitions the best CVC model equals the standard cross-validation model.

6 Discussion

CVC provides a formal framework to take into account the uncertainty in comparing model quality using cross-validation. The crucial difference of CVC from the standard V-fold cross-validation is that it outputs a p-value for each fitted candidate model. The main

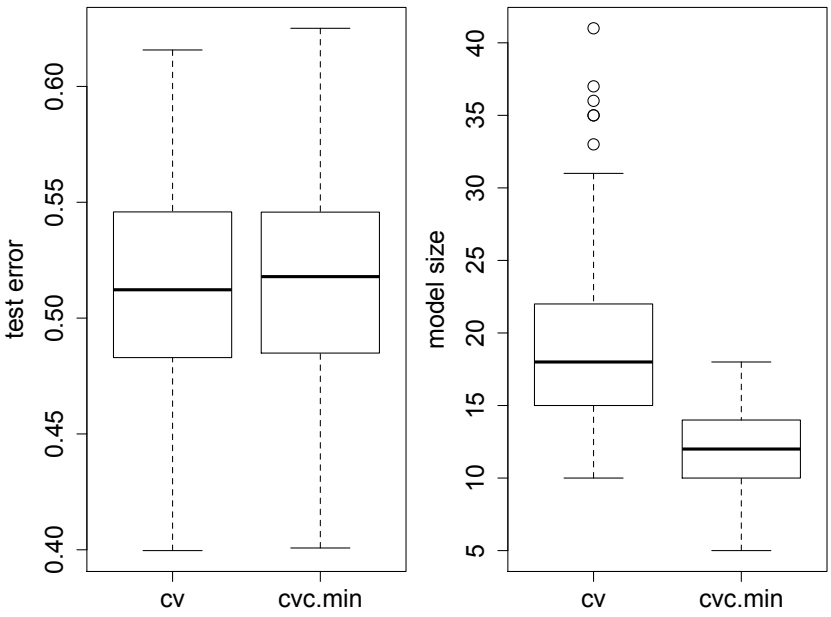

Figure 3: Diabetes data example. Box-plots of hold-out test errors and number of selected variables over 100 repeated hold-out splitting. Compared estimators are the standard cross-validation ("cv") and the most parsimonious model in $\mathcal{A}_{\text{cvc}}$ ("cvc.min").

additional computational cost is the Gaussian multiplier bootstrap used for computing the p-values. [5] described several tricks to speed up the p-value calculation. First, if a model m is clearly inferior than others, then one can directly compare this model with $\hat{m}_{\text{cv}}$ (the model chosen by standard cross-validation) using a paired two sample studentized t-test on their cross-validated prediction errors, calibrated by a Bonferroni correction due to the randomness of $\hat{m}_{\text{cv}}$. If this adjusted two-sample p-value is very small (e.g., smaller than a vanishing number $\alpha'_n \ll \alpha$), then we can skip the bootstrap step and assign $\hat{p}_m = 0$. Secondly, if $\hat{\mu}_{m,m'}/\hat{\sigma}_{m,m'} \ll 0$ for some m and m', then we can confidently infer that $Q(\hat{f}_m) < Q(\hat{f}_{m'})$ and hence remove model m' in the bootstrap calculation of the p-value for model m. This can substantially reduce the complexity of the bootstrap step.

The wide applicability of cross-validation makes it worthwhile to further extend our understanding of CVC under other contexts. In our consistency analysis of sparse linear regression in Section 4.2, we considered independent noise. If the response variable is binary, then we need to extend the theory to cover heteroskedastic noises. It is also possible to extend the framework of CVC to unsupervised learning problems. For example, in K-means or model-based clustering, one needs to specify the number of clusters. Given a set of

estimated clusters, the value of loss function at a test sample point can be set as the squared distance to the closest cluster center. Similar extensions can also be made in low rank matrix approximation using singular value decompositions [16] and network community detection [3], where various forms of cross-validation can be used for model selection.

A Proofs

A.1 Proof of Section 4.1

In the construction of our p-value $\hat{p}_m$, the main test statistic is the maximum of the studentized sample mean of $(\xi_m^{(i)} : i \in I_{\text{te}})$. The reference distribution is the maximum of a Gaussian vector with covariance $\hat{\Gamma}_m$, the sample correlation matrix of $(\xi_m^{(i)} : i \in I_{\text{te}})$. Therefore, the main task is to prove that the maximum of a studentized sample mean has similar distribution as the maximum of a Gaussian random vector with same covariance structure. A result of this type is given as part of the proof of Theorem 4.3 of [5]. Here we provide a self-contained proof using more compact notation and under the assumption of bounded Orlicz ψ_2 norm for $(\xi_{m,j} - \mu_{m,j})/\sigma_{m,j}$. Finally, the unnumbered constants c, c' may vary from line to line. Numbered constants, c_1, c_2, c_3, C_1, C_2, etc, correspond to those specified in corresponding theorems.

Additional notation. Here we fix m and will drop m in the notation. Let μ, Σ, Γ be the mean, covariance and correlation matrix of the $M-1$ dimensional random vector $\xi = (\xi_{m,j} : j \neq m)$ (conditional on D_{tr}). The corresponding sample versions are $\hat{\mu}$, $\hat{\Sigma}$, and $\hat{\Gamma}$. Let σ_j^2 be the jth diagonal element of Σ, and $\hat{\sigma}_j$ the sample version. Finally, let $N = M \vee n$.

Lemma A.1. *Let W_1, W_2 be random variables, then*

$$\|W_1 W_2\|_{\psi_1} \leq \|W_1\|_{\psi_2} \|W_2\|_{\psi_2} .$$

Proof.

$$\begin{aligned} \left\| \frac{W_1 W_2}{\|W_1\|_{\psi_2} \|W_2\|_{\psi_2}} \right\|_{\psi_1} &\leq \left\| \frac{W_1^2}{2\|W_1\|_{\psi_2}^2} + \frac{W_2^2}{2\|W_2\|_{\psi_2}^2} \right\|_{\psi_1} \\ &\leq \frac{1}{2} \left\| \frac{W_1^2}{\|W_1\|_{\psi_2}^2} \right\|_{\psi_1} + \frac{1}{2} \left\| \frac{W_2^2}{\|W_2\|_{\psi_2}^2} \right\|_{\psi_1} = 1 . \end{aligned}$$

■

Lemma A.2. *There exist universal positive constants c_1, c_2 such that if (6) holds with $(A_n^2+1) \le c_1\sqrt{n/\log N}$ then the following holds.*

$$\mathbb{P}\left[\max_j \left|\frac{\hat{\mu}_j - \mu_j}{\sigma_j}\right| \ge c_2 A_n \sqrt{\log N/n}\right] \le N^{-1}\,. \tag{11}$$

$$\mathbb{P}\left[\max_j \left|\frac{\hat{\sigma}_j^2}{\sigma_j^2} - 1\right| \ge c_2 A_n^2 \sqrt{\log N/n}\right] \le N^{-1}\,. \tag{12}$$

$$\mathbb{P}\left[\max_{j,j'} \left|\hat{\Gamma}_{j,j'} - \Gamma_{j,j'}\right| \ge c_2 A_n^2 \sqrt{\log N/n}\right] \le N^{-1}. \tag{13}$$

Proof of Lemma A.2. Since a smaller A_n means better concentration, without loss of generality we can assume that $A_n \ge 1$.

The first inequality directly follows from the fact that $\|Z\|_{\psi_1} \le \|Z\|_{\psi_2}$, the Bernstein's inequality using Orlicz norm (Lemma 2.2.11 of [22]), and union bound.

For the second inequality, note that

$$\left\| \left| \left(\frac{\xi_j - \mu_j}{\sigma_j}\right)^2 - 1 \right|^{1/2} \right\|_{\psi_1} \le \left\| \frac{\xi_j - \mu_j}{\sigma_j} \right\|_{\psi_1} + \|1\|_{\psi_1} = A_n + (\log 2)^{-1}\,. \tag{14}$$

We can re-write $\hat{\sigma}_j^2/\sigma_j^2 - 1$ as

$$\frac{\hat{\sigma}_j^2 - \sigma_j^2}{\sigma_j^2} = n_{\mathrm{te}}^{-1} \sum_{i \in I_{\mathrm{te}}} \left[\left(\frac{\xi_j^{(i)} - \mu_j}{\sigma_j}\right)^2 - 1 \right] - \left(\frac{\hat{\mu}_j - \mu_j}{\sigma_j}\right)^2 .$$

Then the claim follows by applying Cramer's moderate deviation (see, e.g., Theorem 2.13 of [8]) with centered and scaled random variable $A_n^{-2}[(\xi_j - \mu_j)^2/\sigma_j^2 - 1]$ and deviation $c\sqrt{\log N}$, together with (11) and union bound.

For the third inequality, we can re-write the off-diagonal error in Σ

$$\hat{\Sigma}_{j,j'} - \Sigma_{j,j'} = n_{\mathrm{te}}^{-1} \sum_{i \in I_{\mathrm{te}}} (\xi_j^{(i)} - \mu_j)(\xi_{j'}^{(i)} - \mu_{j'}) - (\hat{\mu}_j - \mu_j)(\hat{\mu}_{j'} - \mu_{j'})$$

A similar argument as (14) shows that

$$\left\| \left| \frac{\xi_j - \mu_j}{\sigma_j} \frac{\xi_{j'} - \mu_{j'}}{\sigma_{j'}} - \Gamma_{j,j'} \right|^{1/2} \right\|_{\psi_1} \le A_n + (\log 2)^{-1}\,.$$

Then the same argument leads to the desired result by combining the deviation bound on $\hat{\Sigma} - \Sigma$ and (12). ∎

Proof of Theorem 4.1. Let E denote the event that $\max_j |\hat{\mu}_j - \mu_j|/\sigma_j \le c_2 A_n \sqrt{\log N/n}$, $\max_j |\hat{\sigma}_j/\sigma_j - 1| \le c_2 A_n^2 \sqrt{\log N/n}$, and $\max_{j,j'} |\hat{\Gamma}_{j,j'} - \Gamma_{j,j'}| \le c_2 A_n^2 \sqrt{\log N/n}$. Lemma A.2 implies that under the condition of theorem we have $\mathbb{P}(E) \ge 1 - 3N^{-1}$.

Let Λ and $\hat{\Lambda}$ be the diagonal matrix corresponding to Σ and $\hat{\Sigma}$, respectively. For any positive semidefinite matrix Σ, let Z_Σ be an $N(0, \Sigma)$ random vector. For $\alpha \in (0,1)$, let $z(\alpha, \Sigma)$ be the upper α quantile of the maximum of Z_Σ.

For the first part of the theorem, by assumption we have $\max_j \mu_j/\sigma_j \le \gamma_n$ with $\gamma_n = r_n\sqrt{1/(n \log N)}$. Recall that $\hat{p}$ is the bootstrap p-value given in Section 3.1. Here we ignore the bootstrap sampling variability and regard $\hat{p}$ as the limiting value when the bootstrap sample $B \to \infty$. We have

$$
\begin{aligned}
&\mathbb{P}(\hat{p} \le \alpha) \\
=&\mathbb{P}\left[\sqrt{n_{\mathrm{te}}} \max_j \left(\frac{\hat{\mu}_j - \mu_j}{\hat{\sigma}_j} + \frac{\mu_j}{\hat{\sigma}_j}\right) \ge z(\alpha, \hat{\Gamma})\right] \\
\le&\mathbb{P}\left[\sqrt{n_{\mathrm{te}}} \max_j \left(\frac{\hat{\mu}_j - \mu_j}{\hat{\sigma}_j} + \frac{\mu_j}{\hat{\sigma}_j}\right) \ge z(\alpha, \hat{\Gamma}),\ E\right] + \mathbb{P}(E^c) \\
\le&\mathbb{P}\left[\sqrt{n_{\mathrm{te}}} \max_j \frac{\hat{\mu}_j - \mu_j}{\sigma_j} \ge z(\alpha, \hat{\Gamma}) - c\sqrt{n}\gamma_n - c_2^2 A_n^3 \frac{\log N}{\sqrt{n}},\ E\right] + \mathbb{P}(E^c) \qquad (15) \\
\le&\mathbb{P}\left[\sqrt{n_{\mathrm{te}}} \max_j \frac{\hat{\mu}_j - \mu_j}{\sigma_j} \ge z(\alpha + \delta, \Gamma) - c\sqrt{n}\gamma_n - c_2^2 A_n^3 \frac{\log N}{\sqrt{n}},\ E\right] + \mathbb{P}(E^c) \qquad (16) \\
\le&\mathbb{P}\left[\max Z_\Gamma \ge z(\alpha + \delta, \Gamma) - c\sqrt{n}\gamma_n - c_2^2 A_n^3 \frac{\log N}{\sqrt{n}}\right] + C_1 n^{-C_2} + 3N^{-1} \qquad (17) \\
\le&\alpha + \delta + \left(c\sqrt{n}\gamma_n + c_2^2 A_n^3 \frac{\log N}{\sqrt{n}}\right) 2(1 + \sqrt{2 \log N}) + C_1 n^{-C_2} + 3N^{-1} \qquad (18) \\
\le&\alpha + c \max\left(\left(\frac{A_n^4 \log^5 N}{n}\right)^{\frac{1}{6}}, r_n, \left(\frac{A_n^6 \log^3 N}{n}\right)^{\frac{1}{2}}, N^{-(C_2 \wedge 1)}\right), \qquad (19)
\end{aligned}
$$

where

- (15) follows from the concentration of $\hat{\mu}_j - \mu_j$ and $\hat{\sigma}_j/\sigma_j - 1$, and the definition of E.
- (16) follows from the difference between maxima of Gaussian random vectors (Theorem 2 of [6]). Here $\delta = cA_n^3(\log^{5/6} N)n^{-1/6}$ for some universal constant c.
- (17) follows from the Gaussian comparison of maxima of mean vectors (Corollary 2.1 of [4]) with the ψ_1 norm condition and $\log^7 N \le cn^{1-c'}$ for some constants c, c'. Here C_1 and C_2 in (17) are universal constants involved in the Gaussian comparison result.
- (18) follows from the anti-concentration of maxima of Gaussian random vectors (Theorem 3 of [6]).

When $\mu_j = 0$ for all j, the above argument also goes through in the other direction. In this case we have $\mathbb{P}(\hat{p} \le \alpha) = \alpha + o(1)$. The first part of the theorem is complete.

For the second part. Let j be such that $\frac{\mu_j}{\sigma_j} \ge r\sqrt{\frac{\log N}{n}}$. Then

$$\begin{aligned}\mathbb{P}(\hat{p} \le \alpha) =& \mathbb{P}\left[\sqrt{n_{\text{te}}} \max_{j'} \frac{\hat{\mu}_{j'}}{\hat{\sigma}_{j'}} \ge z(\alpha, \hat{\Gamma})\right] \\ \ge& \mathbb{P}\left[\sqrt{n_{\text{te}}} \frac{\hat{\mu}_j}{\hat{\sigma}_j} \ge z(\alpha, \hat{\Gamma})\right] \\ \ge& \mathbb{P}\left[\sqrt{n_{\text{te}}} \frac{\hat{\mu}_j}{\hat{\sigma}_j} \ge z(\alpha, \hat{\Gamma}), E\right]\end{aligned}$$

On event E, we have

$$\begin{aligned}\frac{\hat{\mu}_j}{\hat{\sigma}_j} &= \left(\frac{\mu_j}{\sigma_j} + \frac{\hat{\mu}_j - \mu_j}{\sigma_j}\right)\left(\frac{\sigma_j}{\hat{\sigma}_j} - 1 + 1\right) \\ &\ge (r - c_2 A_n)\sqrt{\frac{\log N}{n}}\left(1 - c_2 A_n^2 \sqrt{\frac{\log N}{n}}\right) \ge c\sqrt{\log N/n}\,,\end{aligned}$$

for some constant $c > 0$ when n is large enough. Here the last inequality follows because by assumption we have $r \ge c'(A_n \vee 1)$ for some large enough constant c'.

On the other hand, using Mill's inequality and union bound we can directly verify that $z(\alpha, \hat{\Gamma}) \le \sqrt{2(\log(M-1) - \log\alpha)} \le 2\sqrt{\log N}$ whenever $\alpha \ge n^{-1}$. Thus when n is large enough we always have $\sqrt{n_{\text{te}}}\hat{\mu}_j/\hat{\sigma}_j \ge z(\alpha, \hat{\Gamma})$ on E. As a result,

$$\mathbb{P}(\hat{p} \le \alpha) \ge \mathbb{P}(E) \ge 1 - 3N^{-1}.$$

■

A.2 Proofs for Section 4.2

Proof of Theorem 4.2. Let $\mathcal{M} = \{1, 2, ..., M\}$, $\mathcal{M}_1 = \{m \in \mathcal{M} : J_{m^*} \backslash J_m \ne \emptyset\}$ be the candidate models missing at least one true variable, and $\mathcal{M}_2 = \mathcal{M} \backslash (\mathcal{M}_1 \cup \{m^*\})$. Let $\mathcal{A}_n$ be the admissible set given by CVC. We use "$a \lesssim b$" to denote that $a \le cb$ for some universal constant c.

For any $m \in \mathcal{M}$, let J_m be the subset of active variables in model m. Then $\sup_{m\in\mathcal{M}} \|\hat{\beta}_m - \beta_m\|_2 = O_{P_{\text{tr}}}(1/\sqrt{n})$, where $\beta_m = \Sigma^{-1}_{J_m,J_m}\mathbb{E}(X^T_{J_m}Y)$ is the population version, and the notation $O_{P_{\text{tr}}}(\cdot)$ emphasizes that the randomness is on the training data D_{tr}. In the following we overload the notation of $\hat{\beta}_m$ and β_m by embedding them in $\mathbb{R}^p$, with zero values in the entries in J_m^c.

Let E_1 denote the event that $\sup_m \|\hat{\beta}_m - \beta_m\|_2 \le \sqrt{\log n/n}$. The assumptions imply that $\mathbb{P}(E_1) = 1 - o(1)$. Our analysis will condition on E_1.

For two candidate models m, m', the difference of squared error is

$$\begin{aligned}\xi_{m,m'} =&(Y - X^T\hat{\beta}_m)^2 - (Y - X^T\hat{\beta}_{m'})^2 \\ =&2\epsilon X^T(\hat{\beta}_{m'} - \hat{\beta}_m) + (X^T(\hat{\beta}_m - \beta))^2 - (X^T(\hat{\beta}_{m'} - \beta))^2 \,.\end{aligned}$$

We first bound the ψ_1 norm of $\xi_{m,m'}$ for all $m \neq m' \in \mathcal{M}$. First, using Lemma A.1 we have, on E_1,

$$\begin{aligned}&\|\xi_{m,m'}\|_{\psi_1} \\ \lesssim&\|\epsilon\|_{\psi_2}\|X^T(\hat{\beta}_m - \hat{\beta}_{m'})\|_{\psi_2} + \|X^T(\hat{\beta}_m + \hat{\beta}_{m'} - 2\beta)\|_{\psi_2}\|X^T(\hat{\beta}_m - \hat{\beta}_{m'})\|_{\psi_2} \\ \lesssim&\|\epsilon\|_{\psi_2}\|\tilde{X}\|_{\psi_2}\|\Sigma^{1/2}(\hat{\beta}_m - \hat{\beta}_{m'})\|_2 + \|\tilde{X}\|_{\psi_2}^2\beta_{\max}\sqrt{\lambda_{\max}}\|\Sigma^{1/2}(\hat{\beta}_m - \hat{\beta}_{m'})\|_2 \,, \end{aligned} \tag{20}$$

and

$$\begin{aligned}\left|\mathbb{E}(X_{m,m'} \mid D_{\text{tr}})\right| =& \left|(\hat{\beta}_m + \hat{\beta}_{m'} - 2\beta)^T\Sigma(\hat{\beta}_m - \hat{\beta}_{m'})\right| \\ \lesssim&\beta_{\max}\sqrt{\lambda_{\max}}\|\Sigma^{1/2}(\hat{\beta}_m - \hat{\beta}_{m'})\|_2 \,.\end{aligned} \tag{21}$$

On the other hand

$$\text{Var}(\xi_{m,m'} \mid D_{\text{tr}}) \geq \text{Var}(2\epsilon X^T(\hat{\beta}_m - \hat{\beta}_{m'})) = 4\sigma^2\|\Sigma^{1/2}(\hat{\beta}_m - \hat{\beta}_{m'})\|_2^2 \,. \tag{22}$$

Combining (20), (21), and (22) we have

$$\left\|\frac{\xi_{m,m'} - \mathbb{E}(\xi_{m,m'} \mid D_{\text{tr}})}{\sqrt{\text{Var}(\xi_{m,m'} \mid D_{\text{tr}})}}\right\|_{\psi_1} \lesssim \frac{\|\epsilon\|_{\psi_2}\|\tilde{X}\|_{\psi_2} + \|\tilde{X}\|_{\psi_2}^2\beta_{\max}\sqrt{\lambda_{\max}} + \beta_{\max}\sqrt{\lambda_{\max}}}{\sigma} \,, \tag{23}$$

which provides a uniform upper bound that is independent of n and uniform over all m, m'.

Now we show that for any $m \in \mathcal{M}_1$, we have $P(m \in \mathcal{A}_n) \to 0$.

Using the first equation in (21) we have, on E_1

$$\beta_{\min}^2\lambda_{\min} \lesssim \mathbb{E}(\xi_{m,m^*}|D_{\text{tr}}) \lesssim \beta_{\max}^2\lambda_{\max} \,. \tag{24}$$

and

$$\begin{aligned}&\text{Var}(\xi_{m,m^*} \mid D_{\text{tr}}) \\ =&4\sigma^2(\hat{\beta}_m - \hat{\beta}_{m^*})^T\Sigma(\hat{\beta}_m - \hat{\beta}_{m^*}) + \mathbb{E}\left[(\hat{\beta}_m + \hat{\beta}_{m^*} - 2\beta)^T(XX^T - \Sigma)(\hat{\beta}_m - \hat{\beta}_{m^*})\right]^2 \\ \lesssim&\sigma^2\lambda_{\max}\beta_{\max}^2 + \beta_{\max}^4\lambda_{\max}^2 \,.\end{aligned} \tag{25}$$

Combining (24) and (25) we have

$$\frac{\mathbb{E}(\xi_{m,m^*} \mid D_{\text{tr}})}{\sqrt{\text{Var}(\xi_{m,m^*} \mid D_{\text{tr}})}} \geq \frac{\sqrt{\lambda_{\min}}\beta_{\min}}{\sigma\sqrt{\lambda_{\max}}\beta_{\max} + \beta_{\max}^2\lambda_{\max}} \,. \tag{26}$$

(26) and (23) enables us to apply the second part of Theorem 4.1. So we conclude that

$$\mathbb{P}(m \in \mathcal{A}_{\text{cvc}}) \leq P(E_1^c) + \alpha_n + o(1) = o(1) \,. \tag{27}$$

Now consider $m \in \mathcal{M}_2$. For brevity we denote J_m as J and J_{m^*} as J^*. Let $\vec{\epsilon}$ denote the realized vector of ϵ in D_{tr}, X_J denote the *training data* columns of X in J, and Σ_J denote the corresponding population covariance matrix.

$$\begin{aligned}
&\mathbb{E}\left(\left[X^T(\hat{\beta}_m - \beta)\right]^2 | D_{\text{tr}}\right) \\
=&\vec{\epsilon}^T X_J (X_J^T X_J)^{-1} \Sigma_J (X_{J,\text{tr}}^T X_J)^{-1} X_J^T \vec{\epsilon} \\
=&\vec{\epsilon}^T X_J (X_J^T X_J)^{-1} \left[\Sigma_J - (X_J^T X_J/n_{\text{tr}})\right] (X_J^T X_J)^{-1} X_J^T \vec{\epsilon} \\
&\quad + n_{\text{tr}}^{-1} \vec{\epsilon}^T X_J (X_J^T X_J)^{-1} (X_J^T X_J) (X_J^T X_J)^{-1} X_J^T \vec{\epsilon} \\
=&\vec{\epsilon}^T X_J (X_J^T X_J)^{-1} \left[\Sigma_J - (X_J^T X_J/n_{\text{tr}})\right] (X_J^T X_J)^{-1} X_J^T \vec{\epsilon} \\
&\quad + n_{\text{tr}}^{-1} \vec{\epsilon}^T X_J (X_J^T X_J)^{-1} X_J^T \vec{\epsilon} \\
=&O_P(n^{-3/2}) + n_{\text{tr}}^{-1} \vec{\epsilon}^T P_J \vec{\epsilon},
\end{aligned}$$

where P_J is the projection matrix of the linear subspace spanned by the columns of X_J. Therefore

$$n_{\text{tr}}\mathbb{E}(\xi_{m,m^*} \mid D_{\text{tr}}) = \vec{\epsilon}^T (P_J - P_{J^*}) \vec{\epsilon} + O_p(n^{-1/2}).$$

Since $P_J - P_{J^*}$ is a projection matrix, we can write $P_J - P_{J^*} = \sum_{l=1}^{L} v_l v_l^T$, where $L \geq 1$ is the number of irrelevant variables in model J, and $v_1, ..., v_L$ are orthonormal vectors in $\mathbb{R}^{n_{\text{tr}}}$. Thus $\vec{\epsilon}^T (P_J - P_{J^*}) \vec{\epsilon} \geq (v_1^T \vec{\epsilon})^2$.

Given constants (c, c'), let $E_2(c, c') = \{\|X\|_\infty \leq c\sqrt{\log n},\ \lambda_{\min}(X_J^T X_J) \geq c' n\}$, where $\|\cdot\|_\infty$ denotes the largest absolute value. A simple application of concentration of sample mean and union bound implies that, there exists constants c, c', depending only on the ψ_2 norm of X, such that $\mathbb{P}(E_2) \geq 1 - o(1)$. On this event E_2, we have $\|v_1\|_\infty \leq C\sqrt{\log n/n}$, with C determined by (c, c').

Because ϵ has mean zero and finite variance, there exist constants $L > 0$ and $\eta \in (0, 1)$, such that $\sup_t P(\epsilon \in [t, t + L]) \leq 1 - \eta$. Let $v_1 = (v_{1,j} : 1 \leq j \leq n)$. If $v_{1,j} \neq 0$, then $\sup_t P(v_{1,j}\epsilon_j \in [t, t + |v_{1,j}|L]) \leq 1 - \eta$. Now applying the anti-concentration inequality for sums of independent random variables (Theorem 1 of [17]), we have

$$\mathbb{P}\left((v_1^T \vec{\epsilon})^2 \leq n^{-2/5}\right) = \mathbb{P}\left(|v_1^T \vec{\epsilon}| \leq n^{-1/5}\right)$$

$$\lesssim \frac{n^{-1/5}}{\sqrt{\sum_{j:v_{1,j}\neq 0} v_{1,j}^2 L^2(1-\eta)}} \lesssim n^{-1/5}\,.$$

Note that $n^{-1/2}/n^{-2/5} \to 0$, and $(v_1^T\vec{\epsilon})^2 > n^{-2/5}$ implies that $\mathbb{E}(\xi_{m,m^*}) > 0$. Therefore, with probability $1 - o(1)$, $\mathbb{E}(\xi_{m,m^*}) > 0$ for all $m \in \mathcal{M}_2$.

The claimed result follows from Theorem 4.1. ■

References

[1] David M Allen. The relationship between variable selection and data agumentation and a method for prediction. *Technometrics*, 16(1):125–127, 1974.

[2] Sourav Chatterjee and Jafar Jafarov. Prediction error of cross-validated lasso. *arXiv preprint arXiv:1502.06291*, 2015.

[3] Kehui Chen and Jing Lei. Network cross-validation for determining the number of communities in network data. *Journal of the American Statistical Association*, (just-accepted), 2016.

[4] Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41(6):2786–2819, 2013.

[5] Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Testing many moment inequalities. *arXiv preprint arXiv:1312.7614*, 2013.

[6] Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Comparison and anti-concentration bounds for maxima of gaussian random vectors. *Probability Theory and Related Fields*, 162(1-2):47–70, 2015.

[7] Denis Chetverikov and Zhipeng Liao. On cross-validated lasso. *arXiv preprint arXiv:1605.02214*, 2016.

[8] Victor H de la Peña, Tze Leung Lai, and Qi-Man Shao. *Self-normalized processes: Limit theory and Statistical Applications*. Springer Science & Business Media, 2008.

[9] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.

[10] Yang Feng and Yi Yu. Consistent cross-validation for tuning parameter selection in high-dimensional variable selection. *arXiv preprint arXiv:1308.5390*, 2013.

[11] Seymour Geisser. The predictive sample reuse method with applications. *Journal of the American Statistical Association*, 70(350):320–328, 1975.

[12] Darren Homrighausen and Daniel J McDonald. Risk-consistency of cross-validation with lasso-type procedures. *arXiv preprint arXiv:1308.0810*, 2013.

[13] Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, to appear, 2017.

[14] Ker-Chau Li. Asymptotic optimality for cp, cl, cross-validation and generalized cross-validation: discrete index set. *The Annals of Statistics*, pages 958–975, 1987.

[15] Chinghway Lim and Bin Yu. Estimation stability with cross-validation (escv). *Journal of Computational and Graphical Statistics*, 25(2):464–492, 2016.

[16] Art B Owen and Patrick O Perry. Bi-cross-validation of the svd and the nonnegative matrix factorization. *The annals of applied statistics*, pages 564–594, 2009.

[17] Boris Alekseevich Rogozin. On the increase of dispersion of sums of independent random variables. *Theory of Probability & Its Applications*, 6(1):97–99, 1961.

[18] Jun Shao. Linear model selection by cross-validation. *Journal of the American statistical Association*, 88(422):486–494, 1993.

[19] Mervyn Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 111–147, 1974.

[20] Ryan J Tibshirani and Robert Tibshirani. A bias correction for the minimum error rate in cross-validation. *The Annals of Applied Statistics*, pages 822–829, 2009.

[21] Sara van de Geer and Johannes Lederer. The bernstein–orlicz norm and deviation inequalities. *Probability Theory and Related Fields*, 157(1-2):225–250, 2013.

[22] Aad W van der Vaart and Jon A Wellner. *Weak Convergence and Empirical Processes*. Springer, 1996.

[23] Larry Wasserman and Kathryn Roeder. High dimensional variable selection. *Annals of statistics*, 37(5A):2178, 2009.

[24] Yuhong Yang. Consistency of cross validation for comparing regression procedures. *The Annals of Statistics*, pages 2450–2473, 2007.

[25] Ping Zhang. Model selection via multifold cross validation. *The Annals of Statistics*, pages 299–313, 1993.